RESEARCH

Near-optimal Assembly for Shotgun Sequencing with Noisy Reads

Ka-Kit Lam[1], Asif Khalak[2] and David Tse[1*]

Abstract

Recent work identified the fundamental limits on the information requirements in terms of read length and coverage depth required for successful *de novo* genome reconstruction from shotgun sequencing data, based on the idealistic assumption of no errors in the reads (noiseless reads). In this work, we show that even when there is noise in the reads, one can successfully reconstruct with information requirements close to the noiseless fundamental limit. A new assembly algorithm, X-phased Multibridging, is designed based on a probabilistic model of the genome. It is shown through analysis to perform well on the model, and through simulations to perform well on real genomes.

Keywords: De novo sequence assembly, genome finishing, methods for emerging sequencing technologies

*Correspondence:
dnctse@gmail.com
[1]Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, California, United States
[2] Pacific Biosciences, Menlo Park, California, United States
Full list of author information is available at the end of the article

arXiv:1402.6971v1 [q-bio.GN] 27 Feb 2014

Background

Optimality in the acquisition and processing of DNA sequence data represents a serious technology challenge from various perspectives including sample preparation, instrumentation and algorithm development. Despite scientific achievements such as the sequencing of the human genome and ambitious plans for the future [17, 14], there is no single, overarching framework to identify the fundamental limits in terms of information requirements required for successful output of the genome from the sequence data.

Information theory has been successful in providing the foundation for such a framework in digital communication [15], and we believe that it can also provide insights into understanding the essential aspects of DNA sequencing. A first step in this direction has been taken in the recent work [1], where the fundamental limits on the minimum read length and coverage depth required for successful assembly are identified in terms of the statistics of various repeat patterns in the genome. Successful assembly is defined as the reconstruction of the underlying genome, i.e. genome finishing [13]. The genome finishing problem is particularly attractive for analysis because it is clearly and unambiguously defined and is arguably the ultimate goal in assembly. There is also a scientific need for finished genomes [10][9]. Until recently, automated genome finishing was beyond reach [4] in all but the simplest of genomes. New advances using ultra-long read single-molecule sequencing, however, have reported successful automated finishing [2, 6]. Even in the case where finished assembly is not possible, the results in [1] provide insights on optimal use of read information since the heart of the problem lies in how one can optimally use the read information to resolve repeats.

Figure 1a gives an example result for the repeat statistics of *E. coli* K12. The x-axis of the plot is the read length and the y-axis is the coverage depth normalized

by the Lander-Waterman depth (number of reads needed to cover the genome [8]). The lower bound identifies the necessary read length and coverage depth required for *any* assembly algorithm to be successful with these repeat statistics. An assembly algorithm called Multibridging Algorithm was presented, whose read length and coverage depth requirements are very close to the lower bound, thus tightly characterizing the fundamental information requirements. The result shows a critical phenomenon at a certain read length $L = \ell_{\text{crit}}$: below this critical read length, reconstruction is impossible no matter how high the coverage depth; slightly above this read length, reconstruction is possible with Lander-Waterman coverage depth. This critical read length is given by $\ell_{\text{crit}} = \max\{\ell_{\text{int}}, \ell_{\text{tri}}\}$, where ℓ_{int} is the length of the longest pair of exact interleaved repeats and ℓ_{tri} is the length of the longest exact triple repeat in the genome, and has its roots in earlier work by Ukkonen on Sequencing-by-Hybridization [18]. The framework also allows the analysis of specific algorithms and the comparison with the fundamental limit; the plot shows for example the performance of the Greedy Algorithm and we see that its information requirement is far from the fundamental limit.

A key simplifying assumption in [1] is that there are no errors in the reads (noiseless reads). However reads are noisy in all present-day sequencing technologies, ranging from primarily substitution errors in Illumina® platforms, to primarily insertion-deletion errors in Ion Torrent® and PacBio® platforms. The following question is the focus of the current paper: in the presence of read noise, can we still successfully assemble with a read length and coverage depth close to the minimum in the noiseless case? A recent work [5] with an existing assembler suggests that the information requirement for genome finishing substantially exceeds the noiseless limit. However, it is not obvious whether the limitations lie in the fundamental effect of read noise or in the sub-optimality of the algorithms in the assembly pipeline.

Results

The difficulty of the assembly problem depends crucially on the genome repeat statistics. Our approach to answering the question of the fundamental effect of read noise is based on design and analysis using a parametric probabilistic model of the genome that matches the key features of the repeat statistics we observe in genomes. In particular, it models the presence of long flanked repeats which are repeats flanked by statistically uncorrelated region. Figure 1b shows a plot of the predicted information requirement for reliable reconstruction by various algorithms under a substitution error rate of 1%. The plot is based on analytical formulas derived under our genome model with parameters set to match the statistics of *E. coli* K12. We show that it is possible in many cases to develop algorithms that approach the noiseless lower bound even when the reads are noisy. Specifically, the X-phased Multibridging Algorithm has close to the same critical read length $L = \ell_{\text{crit}}$ as in the noiseless case and only slightly greater coverage depth requirement for read lengths greater than the critical read length.

We then proceed to build a prototype assembler based on the analytical insights and we perform experiments on real genomes. As shown in Figure 2, we test the prototype assembler by using it to assemble noisy reads sampled from 4 different genomes. At coverage and read length indicated by a green circle, we successfully

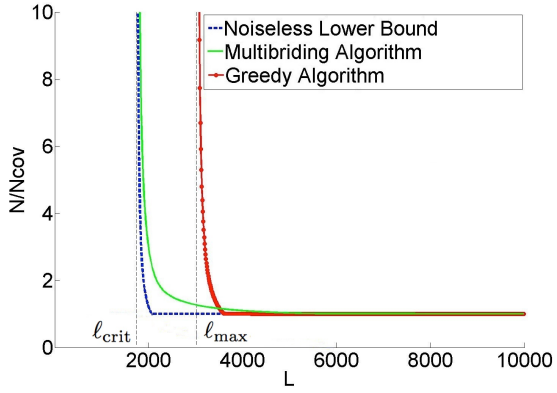

(a) Information requirement for noiseless reads

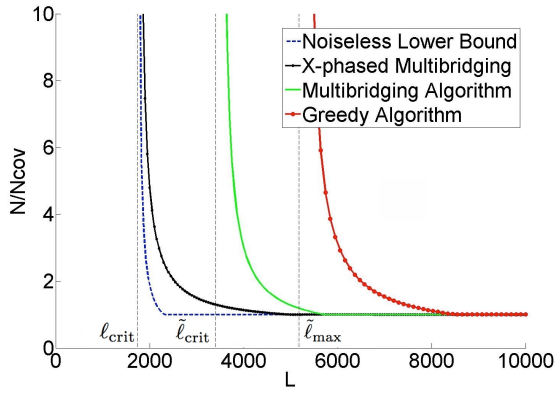

(b) Information requirement for noisy reads

Figure 1: Information requirement to reconstruct *E. coli* K12. $\ell_{\text{crit}} = 1744, \tilde{\ell}_{\text{crit}} = 3393$

assemble noisy reads into one contig (in most cases with more than 99% of the content matched when compared with the ground truth). Note that the information requirement is close to the noiseless lower bound. Moreover, the algorithm (X-phased Multibridging) is computationally efficient with the most computational expensive step being the computation of overlap of reads/K-mers, which is an unavoidable procedure in most assembly algorithms.

The main conclusion of this work is that, with an appropriately designed assembly algorithm, the information requirement for genome assembly is surprisingly insensitive to read noise. The basic reason is that the redundancy required by the Lander-Waterman coverage constraint can be used to denoise the data. This is consistent with the asymptotic result obtained in [11] and the practical approach taken in [2]. However, the result in [11] is based on a very simplistic i.i.d. random genome model, while the model and genomes considered in the present paper both have long repeats. A natural extension of the Multibridging Algorithm in [1] to handle noisy reads allows the resolution of these long flanked repeats if the reads are long enough to span them, thus allowing reconstruction provided that the read length

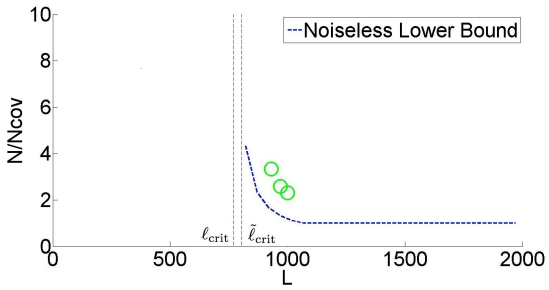

(a) *Prochlorococcus marinus*

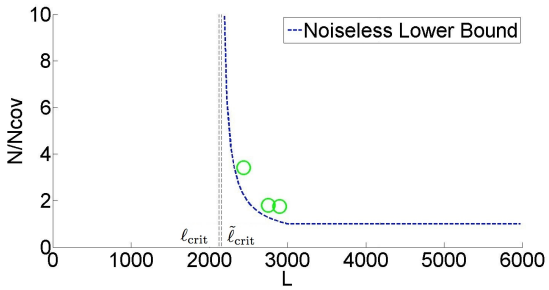

(b) *Helicobacter pylori*

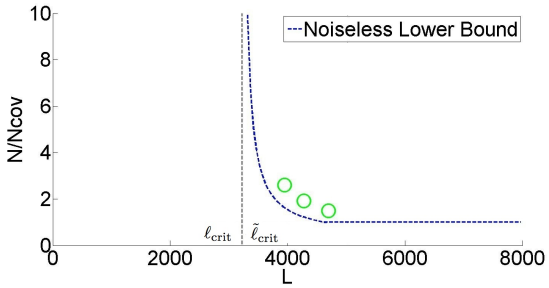

(c) *Methanococcus maripaludis*

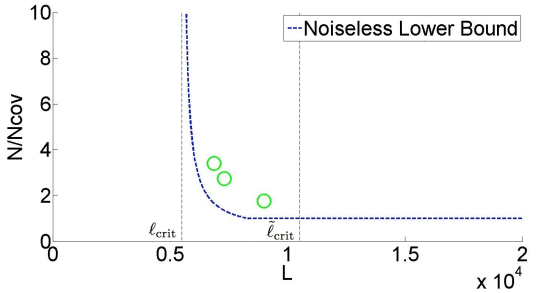

(d) *Mycoplasma agalactiae*

Figure 2: Simulation results on a prototype assembler (substitution noise of rate 1.5 %)

is greater than $L = \tilde{\ell}_{\text{crit}} = \max\{\tilde{\ell}_{\text{int}}, \tilde{\ell}_{\text{tri}}\}$, where $\tilde{\ell}_{\text{int}}$ is the length of the longest pair of flanked interleaved repeats and $\tilde{\ell}_{\text{tri}}$ is the length of the longest flanked triple

repeat in the genome. This condition is shown as a vertical asymptote of the "Multibridging Algorithm" curve in Figure 1b. By exploiting the redundancy in the read coverage to resolve read errors, the X-phased Multibridging can phase the polymorphism across the flanked repeat copies using only reads that span the exact repeats. Hence, reconstruction is achievable with a read length close to $L = \ell_{\text{crit}}$, which is the noiseless limit.

Related work

All assemblers must somehow address the problem of resolving noise in the reads during genome reconstruction. However, the traditional approaches to measuring assembly performance makes quantitative comparisons challenging for unfinished genomes [12]. In most cases, the heart of the assembly problem lies in processing of the assembly graph, as in [19, 3, 16]. A common strategy for dealing with ambiguity from the reads lies in filtering the massively parallel sequencing data using the graph structure prior to traversing possible assembly solutions. In the present work, however, we are focused on the often-overlooked goal of optimal data efficiency. Thus, to the extent possible we distinguish between the read error and the mapping ambiguity associated with the shotgun sampling process. The proposed assembler, X-phased Multibridging, adds information to the assembly graph based on a novel analysis of the underlying reads.

Methods

The path towards developing X-phased Multibridging is outlined as follows.

1. Setting up the shotgun sequencing model and problem formulation.
2. Analyzing repeats structure of genome and their relationship to the information requirement for genome finishing.
3. Developing a parametric probabilistic model that captures the long tail of the repeat statistics.
4. Deriving and analyzing an algorithm that require minimal information requirements for assembly – close to the noiseless lower bound.
5. Performing simulation-based experiments on real and synthetic genomes to characterize the performance of a prototype assembler for genome finishing.
6. Extending the algorithm to address the problem of indel noise.

Shotgun sequencing model and problem formulation

Sequencing model

Let $\mathbf{s}$ be a length G target genome being sequenced with each base in the alphabet set $\Sigma = \{A, C, G, T\}$. In the shotgun sequencing process, the sequencing instrument samples N reads, $\vec{r}_1, \ldots, \vec{r}_N$ of length L and sampled uniformly and independently from $\mathbf{s}$. This unbiased sampling assumption is made for simplicity and is also supported by the characteristics of single-molecule (e.g. PacBio®) data. Each read is a noisy version of the corresponding length L substring on the genome. The noise may consist of base insertions, substitutions or deletions. Our analysis focus on substitution noise first. In a later section, indel noise is addressed. In the substitution noise model, let p be the probability that a base is substituted by another base, with probability $p/3$ to be any other base. The errors are assumed to be independent across bases and across reads.

Formulation

Successful reconstruction by an algorithm is defined by the requirement that, with probability at least $1-\epsilon$, the reconstruction $\hat{\mathbf{s}}$ is a single contig which is within edit distance δ from the target genome $\mathbf{s}$. If an algorithm can achieve that guarantee at some (N, L), it is called ϵ-feasible at (N, L). This formulation implies automated genome finishing, because the output of the algorithm is one single contig. The fundamental limit for the assembly problem is the set of (N, L) for which successful reconstruction is possible by some algorithms. If $\hat{\mathbf{s}}$ is directly spelled out from a correct placement of the reads, the edit distance between $\hat{\mathbf{s}}$ and $\mathbf{s}$ is of the order of pG, where the error rate is p. This motivates fixing $\delta = 2pG$ for concreteness. The quality of the assembly can be further improved if we follow the assembly algorithm with a consensus stage in which we correct each base, e.g. with majority voting. But the consensus stage is not the focus in this paper.

Repeats structure and their relationship to the information requirement for successful reconstruction

Long exact repeats and their relationship to assembly with noiseless reads

We take a moment to carefully define the various types of exact repeats. Let $\mathbf{s}_{\mathbf{t}}^{\ell}$ denote the length-ℓ substring of the DNA sequence $\mathbf{s}$ starting at position t. An exact repeat of length ℓ is a substring appearing twice, at some positions t_1, t_2 (so $\mathbf{s}_{\mathbf{t_1}}^{\ell} = \mathbf{s}_{\mathbf{t_2}}^{\ell}$) that is maximal (i.e. $s(t_1-1) \neq s(t_2-1)$ and $s(t_1+\ell) \neq s(t_2+\ell)$).

Similarly, an exact triple repeat of length-ℓ is a substring appearing three times, at positions t_1, t_2, t_3, such that $\mathbf{s}_{\mathbf{t_1}}^{\ell} = \mathbf{s}_{\mathbf{t_2}}^{\ell} = \mathbf{s}_{\mathbf{t_3}}^{\ell}$, and such that neither of $s(t_1-1) = s(t_2-1) = s(t_3-1)$ nor $s(t_1+\ell) = s(t_2+\ell) = s(t_3+\ell)$ holds.

A copy of a repeat is a single one of the instances of the substring appearances. A pair of exact repeats refers to two exact repeats, each having two copies. A pair of exact repeats, one at positions t_1, t_3 with $t_1 < t_3$ and the second at positions t_2, t_4 with $t_2 < t_4$, is interleaved if $t_1 < t_2 < t_3 < t_4$ or $t_2 < t_1 < t_4 < t_3$. The length of a pair of exact interleaved repeats is defined to be the length of the shorter of the two exact repeats. A typical appearance of a pair of exact interleaved repeat is –X–Y–X–Y– where X and Y represent two different exact repeat copies and the dashes represent non-identical sequence content.

We let $\ell_{\max}$ be the length of the longest exact repeat, ℓ_{int} be the length of the longest pair of exact interleaved repeats and ℓ_{tri} be the length of the longest exact triple repeat.

As mentioned in the introduction, it was observed that the read length and coverage depth required for successful reconstruction using noiseless reads for many genomes is governed by long exact repeats. For some algorithms (e.g. Greedy Algorithm), the read length requirement is bottlenecked by $\ell_{\max}$. The Multibridging Algorithm in [1] can successfully reconstruct the genome with a minimum amount of information. The corresponding minimum read length requirement is the critical exact repeat length $\ell_{\text{crit}} = \max(\ell_{\text{int}}, \ell_{\text{tri}})$.

Flanked repeats

While exact repeats are defined as the segments terminated on each end by a single differing base (Fig 3a), flanked repeats are defined by the segments terminated on

each end by a statistically uncorrelated region. We call that ending region to be the *random flanking region*. A distinguishing characteristic of the random flanking region is a high Hamming distance to segment length ratio between the ends of two repeat copies. The ratio in the random flanking region is around 0.75, which matches with that when the genomic content is independently and uniformly randomly generated. We observe that long repeats of many genomes terminate with random flanking region. Additional statistical analysis is detailed in the Appendix.

If the repeat interior is exactly the same between two copies of the flanked repeat (Fig 3b), the corresponding flanked repeat is called a flanked exact repeat. If there are a few edits (called polymorphism) within the repeat interior (Fig 3c), the corresponding flanked repeat is called a flanked approximate repeat.

The length of the repeat interior bounded by the random flanking region is then the flanked repeat length. We let $\tilde{\ell}_{\max}$ be the length of the longest flanked repeat, $\tilde{\ell}_{\text{int}}$ be the length of the longest pair of flanked interleaved repeats and $\tilde{\ell}_{\text{tri}}$ be the length of the longest flanked triple repeat. The critical flanked repeat length is then $\tilde{\ell}_{\text{crit}} = \max(\tilde{\ell}_{\text{int}}, \tilde{\ell}_{\text{tri}})$.

Long flanked exact repeats and their relationship to assembly with noisy reads

If all long flanked repeats are flanked exact repeats, we can utilize the information in the random flanking region to generalize Greedy Algorithm and Multibridging Algorithm to handle noisy reads. The corresponding information requirement is very similar to that when we are dealing with noiseless reads.

The key intuition is as follows. A criterion for successful reconstruction is the existence of reads to span the repeats to their neighborhood. When a read is noiseless, it only need to be long enough to span the repeat interior to its neighborhood by one base (Fig 4a) so as to differentiate between two exact repeat copies. When a read is noisy, it then need to be long enough to span the repeat interior plus a short extension into the random flanking region (Fig 4b) so as to confidently differentiate between two flanked repeat copies. However, the Hamming distance between two flanked repeat copies' neighborhood in the random flanking region is very high even within a short length. This can be used to differentiate between two flanked repeat copies confidently even when the reads are noisy. The short extension into the random flanking region has a length which is typically of order of tens whereas the long repeat length is of order of thousands. Therefore, relative to the repeat length, the change of the critical read length requirement from handling noiseless reads to noisy reads is only marginal when all long repeats are flanked exact repeats.

Long flanked approximate repeats and their relationship to assembly with noisy reads

If a long flanked repeat is a flanked approximate repeat, the flanked repeat length may be significantly longer than the length of its longest enclosed exact repeat. Merely relying on the information provided by the random flanking region requires the reads to be of length longer than the flanked repeat length for successful reconstruction. This explains why the information requirement for Greedy Algorithm and Multibridging Algorithm has a significant increase when we use noisy reads instead of noiseless reads (as shown in Fig 1b). However, if we utilize the information provided by the coverage, we can still confidently differentiate different repeat

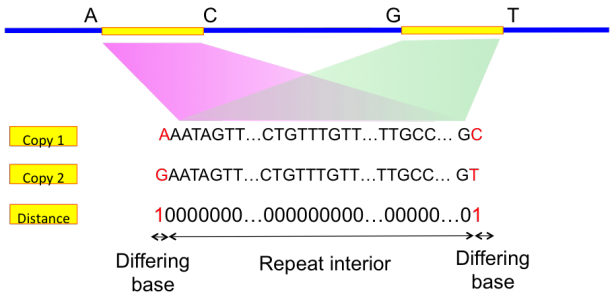

(a) Exact repeat

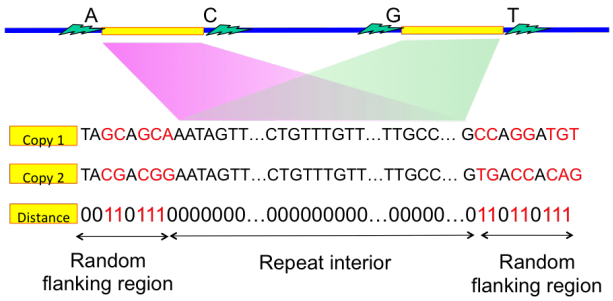

(b) Flanked exact repeat

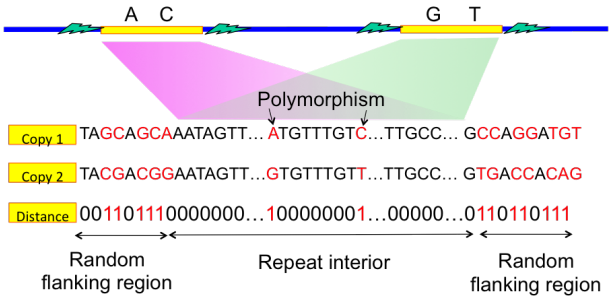

(c) Flanked approximate repeat

Figure 3: Repeat pattern

copies by phasing the small edits within the repeat interior (Fig 4c). Specifically, we design X-phased Multibridging whose information requirement is close to the noiseless lower bound even when some long repeats are flanked approximate repeats, as shown in Fig 1b.

From information theoretic insight to algorithm design

Because of the structure of long flanked repeats, there are two important sources of information that we specifically want to utilize when designing data-efficient algorithms to assemble noisy reads. They are

- The random flanking region beyond the repeat interior
- The coverage given by multiple reads overlapping at the same site

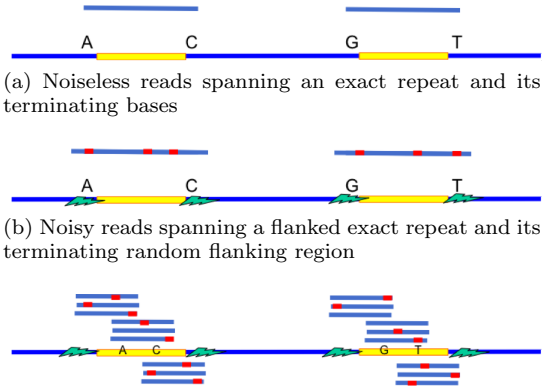

(a) Noiseless reads spanning an exact repeat and its terminating bases

(b) Noisy reads spanning a flanked exact repeat and its terminating random flanking region

(c) Noisy reads extending to span a flanked approximate repeat and its terminating random flanking region

Figure 4: Intuition behind the information requirement

Greedy Algorithm(Alg 1) utilizes the random flanking region when considering overlap. The minimum read length needed for successful reconstruction is close to $\tilde{\ell}_{\max}$.

Multibridging Algorithm(Alg 2) also utilizes the random flanking region but it improves upon Greedy Algorithm by using a De Bruijn graph to aid the resolution of flanked repeats. The minimum read length needed for successful reconstruction is close to $\tilde{\ell}_{\text{crit}}$.

X-phased Multibridging(Alg 3) further utilizes the coverage given by multiple reads to phase the polymorphism within the repeat interior of flanked approximate repeats. The minimum read length needed for successful reconstruction is close to ℓ_{crit}, which is the noiseless lower bound even when some long repeats are flanked approximate repeats.

Model for genome

To capture the key characteristics of repeats and to guide the design of assembly algorithms, we use the following parametric probabilistic model for genome. A target genome is modeled as a random vector $\mathbf{s}$ of length G that has the following three key components (a pictorial representation is depicted in Figure 5).

Random background: The background of the genome is a random vector, composed of uniformly and independently picked bases from the alphabet set $\Sigma = \{A, C, G, T\}$.

Long flanked repeats: On top of the random background, we randomly position the longest flanked repeat and the longest flanked triple repeat. Moreover, we randomly position a flanked repeat interleaving the longest flanked repeat, forming the longest pair of flanked interleaved repeat. The corresponding length of the flanked repeats are $\tilde{\ell}_{\max}$, $\tilde{\ell}_{\text{tri}}$ and $\tilde{\ell}_{\text{int}}$ respectively. It is noted that $\tilde{\ell}_{\max} > \max(\tilde{\ell}_{\text{int}}, \tilde{\ell}_{\text{tri}})$.

Polymorphism and long exact repeats: Within the repeat interior of the flanked repeats, we randomly position $n_{\max}$, n_{int} and n_{tri} edits (polymorphism) respectively. The sites of polymorphism are chosen such that the longest exact

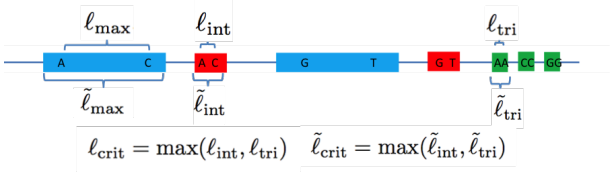

Figure 5: Model for genome

repeat, the longest pair of exact interleaved repeats and the longest exact triple repeat are of length $\ell_{\max}$, ℓ_{int} and ℓ_{tri} respectively.

Algorithm design and analysis

Greedy Algorithm

Read R_2 is a successor of read R_1 if there exists length-W suffix of R_1 and length-W prefix of R_2 such that they are extracted from the same locus on the genome. Furthermore, there is no other reads that can satisfy the same condition with a larger W. To properly determine successors of reads in the presence of long repeats and noise, we need to define an appropriate overlap rule for reads. In this section, we show the conceptual development towards defining such a rule, which is called RA-rule.

Noiseless reads and long exact repeats: If the reads are noiseless, all reads can be paired up with their successors correctly with high probability when the read length exceeds $\ell_{\max}$. It was done [1] by greedily pairing reads and their candidate successors based on their overlap score in descending order. When a read and a candidate successor are paired, they will be removed from the pool for pairing. Here the overlap score between a read and a candidate successor is the maximum length such that the suffix of the read and prefix of the candidate successor match *exactly*.

Noisy reads and random background: Since we cannot expect exact match for noisy reads, we need a different way to define the overlap score. Let us consider the following toy situation. Assume that we have exactly one length-$(\ell+1)$ noisy read starting at each locus of a length G random genome(i.e. only consists of the random background). Each read then overlaps with its successor precisely by ℓ bases. Analogous to the noiseless case, one would expect to pair reads greedily based on overlap score. Here the overlap score between a read and a candidate successor is the maximum length such that the suffix (x) of the read and prefix (y) of the candidate successor match *approximately*. To determine whether they match *approximately*, one can use a predefined a threshold factor α and compute the Hamming distance $d(x, y)$. If $d(x, y) \leq \alpha \cdot \ell$, then they match *approximately*, otherwise not. Given this decision rule, we can have false positive (i.e. having any pairs of reads mistakenly paired up) and false negative (i.e. having any reads not paired up with the true successors). If false positive and false negative probability are small, this naive method is a reliable enough metric. This can be achieved by using a long enough length $\ell > \ell_{\text{iid}}$ and an appropriately chosen threshold α.

Recall that ϵ is the overall failure probability. By bounding the sum of false positive and false negative probability by $\epsilon/3$, one can find $\ell_{\text{iid}}(p, \epsilon/3, G)$ and $\alpha(p, \epsilon/3, G)$ to

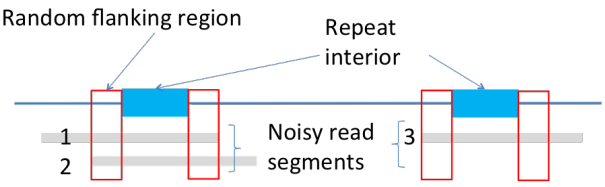

Figure 6: Intuition about why we define the overlap rule to be RA-overlap rule

be the $(\ell_{\text{iid}}, \alpha)$ solution to the following pair of equations:

$$G^2 \cdot \exp(-\ell_{\text{iid}} \cdot D(\alpha||\frac{3}{4})) = \frac{\epsilon}{6} \quad (1)$$

$$G \cdot \exp(-\ell_{\text{iid}} \cdot D(\alpha||2p - \frac{4}{3}p^2)) = \frac{\epsilon}{6} \quad (2)$$

where $D(a||b) = a \log \frac{a}{b} + (1-a) \log \frac{1-a}{1-b}$ is the Kullback-Leibler divergence.

Noisy reads and long flanked repeats: However, when the genome contains long flanked repeats on top of the random background, this naive rule of determining overlap is not enough. Let us look at the example in Fig. 6. As shown in Fig. 6, because of long flanked repeats, we have a small ratio of overall distance against the overlap length for the segments that are extracted from different copies of the repeat (e.g Segment 1 and Segment 3 in Fig. 6). Therefore, the overall Hamming distance between two segments is not a good enough metric for defining overlap. If we abide by the naive rule, we need to increase the read length significantly longer than the flanked repeat length so as to guarantee confidence in deciding approximate match. Otherwise, it will either result in a high false positive rate (if we set a large α) or a high false negative rate (if we set a small α). To properly handle such scenario, we define a repeat-aware rule(or RA-rule).

- RA-matching: Two segments (x, y) of length W match under the RA-rule if and only if the distance between whole segments is $< \alpha \cdot W$ and **both** of its ending segments(of length ℓ_{iid}) also have distance $< \alpha \cdot \ell_{\text{iid}}$.
- RA-overlap: The overlap score between a read and a candidate successor under the RA-rule is the maximum length such that the suffix of the read and prefix of the candidate successor match under the RA-matching.

The RA-rule is particularly useful because it puts an emphasis on both ends of the overlap region. Since the ends are separated by a long range, one end will hopefully originate from the random flanking region of the flanked repeat. If we focus on the segments originating from the random flanking region, the distance per segment length ratio will be very high when the segments originate from different copies of the repeat but very low when they originate from the same copy of the repeat. This is how we utilize the random flanking region to differentiate between repeat copies and determine correct successors in the presence of long flanked repeats and noise.

If we use Greedy Algorithm (Alg 1) to merge reads greedily with this overlap rule (RA-rule), Prop 1 shows the information requirement under the previously

described sequencing model and genome model. A plot is shown in Fig 1b. Since ℓ_{iid} is of order of tens whereas $\tilde{\ell}_{\max}$ is of order of thousands, the read length requirement for Greedy Algorithm to succeed is dominated by $\tilde{\ell}_{\max}$. The detailed proof of Prop 1 is given in Appendix.

Algorithm 1 Greedy Algorithm

```
Initialize contigs to be reads
for W = L to ℓ_iid do
    if any two contigs x,y are of overlap W under RA-rule then
        merge x, y into one contig.
    end
end
```

Proposition 1 *With* $\ell_{\text{iid}} = \ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$, *if*

$$L > \tilde{\ell}_{\max} + 2\ell_{\text{iid}},$$

$$N > \max\left(\frac{G\ln(3/\epsilon)}{L - \tilde{\ell}_{\max} - 2\ell_{\text{iid}}}, \frac{G\ln(3N/\epsilon)}{L - 2\ell_{\text{iid}}}\right)$$

then, Greedy Algorithm(Alg 1) is $\epsilon-$*feasible at* (N, L).

Multibridging Algorithm

The read length requirement of Greedy Algorithm has a bottleneck around $\tilde{\ell}_{\max}$ because it requires at least one copy of each flanked repeat to be spanned by at least one read for successful reconstruction. Spanning a repeat by a single read is called bridging in [1]. A natural question is whether we need to have all repeats bridged for successful reconstruction.

In the noiseless setting, [1] shows that this condition can be relaxed. Using noiseless reads, one can have successful reconstruction given all copies of each exact triple repeat being bridged, and at least one copy of one of the repeats in each pair of exact interleaved repeats being bridged.

A key idea to allow such a relaxation in [1] is to use a De Bruijn graph to capture the structure of the genome.

When the reads are noisy, we can utilize the random flanking region to specify a De Bruijn graph with high confidence by RA-rule and arrive at a similar relaxation. By some graph operations to handle the residual errors, we can have successful reconstruction with read length $\tilde{\ell}_{\text{crit}} + 2 \cdot \ell_{\text{iid}} < L < \tilde{\ell}_{\max}$. The algorithm is summarized in Alg 2. Prop 2 shows its information requirement under the previously described sequencing model and genome model. A plot is shown in Fig 1b. We note that Alg 2 can be seen as a noisy reads generalization of Multibridging Algorithm for noiseless reads in [1].

Description and its performance

Proposition 2 *With* $\ell_{\text{iid}} = \ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$, if

$$L > \tilde{\ell}_{\text{crit}} + 2\ell_{\text{iid}},$$

$$N \;>\; \max\left(\frac{G\ln(3/\epsilon)}{L-\tilde{\ell}_{\text{crit}}-2\ell_{\text{iid}}},\; \frac{G\ln(3N/\epsilon)}{L-2\ell_{\text{iid}}}\right)$$

then, Multibridging Algorithm(Alg 2) is $\epsilon-$feasible at (N, L).

Algorithm 2 Multibridging Algorithm

1. Choose K to be $\tilde{\ell}_{\text{crit}}+2\ell_{\text{iid}}$ and extract K-mers from reads.
2. Cluster K-mers based on the RA-rule.
3. Form uncondensed De Bruijn graph $G_{De-Bruijn}=(V,E)$ with the following rule:
   - a) K-mers clusters as node set V.
   - b) $(u,v)=e\in E$ if and only if there exists K-mers $u_1\in u$ and $v_1\in v$ such that u_1,$v_1$are consecutive K-mers in some reads.
4. Join the disconnected components of $G_{De-Bruijn}$ together by the following rule:

for $W=K-1$ *to* ℓ_{iid} **do**
  for *each node u which has either no predecessors / successors in $G_{De-Bruijn}$* **do**
    a) Find the predecessor/successor v for u from all possible K-mers clusters such that overlap length(using any representative K-mers in that cluster) between u and v is W under RA-rule.
    b) Add dummy nodes in the De Bruijn graph to link u with v and update the graph to $G_{De-Bruijn}$
  end
end

5. Condense the graph $G_{De-Bruijn}$ to form G_{string} with the following rule:
   - a) Initialize G_{string} to be $G_{De-Bruijn}$ with node labels of each node being its cluster group index.
   - b) **while** $\exists$*successive nodes $u\to v$ such that $out-degree(u)=1 and\ in-degree(v)=1$* **do**
     bi) Merge u and v to form a new node w
     bii) Update the node label of w to be the concatenation of node labels of u and v
     end

6. Clear Branches of G_{string}:

for *each node u in the condensed graph G_{string}* **do**
  if *$out-degree(u)>1$ and that all the successive paths are of the same length(measured by the number of node labels) and then joining back to node v and the path length $<\ell_{\text{iid}}$* **then**
    we merge the paths into a single path from u to v.
  end
end

7. Condense graph G_{string}
8. Find the genome :
   - a) Find an Euler Cycle/Path in G_{string} and output the concatenation of the node labels to form a string $\vec{s}_{labels}$.
   - b) Using $\vec{s}_{labels}$ and look up the associated K-mers to form the final recovered genome $\hat{s}$.

Detailed proof is given in the Appendix. The following sketch highlights the motivation behind the key steps of Multibridging Algorithm.

[Step1] We set a large K value to make sure the K-mers overlapping the shorter repeat of the longest pair of flanked interleaved repeats and the longest flanked triple repeat can be separated as distinct clusters.

[Step2] Clustering is done using the RA-rule because of the existence of long flanked repeats and noise.

[Step3] A K-mer cluster corresponds to an equivalence class for K-mers matched under the RA-rule. This step forms a De Bruijn graph with K-mer clusters as nodes.

[Step4] Because of large K, the graph can be disconnected due to insufficient coverage. In order to reduce the coverage constraint, we connect the clusters greedily.

[Step5, 7] These two steps simplify the graph.

[Step6] Branch clearing repairs any incorrect merges near the boundary of long flanked repeat.

[Step8] Since an Euler path in the condensed graph corresponds to the correct genome sequence, it is traversed to form the reconstructed genome.

Some implementation details: improvement on time and space efficiency

For Multibridging Algorithm, the most computational expensive step is the clustering of K-mers. To improve the time and space efficiency, this clustering step can be approximated by performing pairwise comparison of reads.

Based on the alignment of the reads, we can cluster K-mers from different reads together using a disjoint set data structure that supports union and find operations. Since only reads are used in the alignment, only the K-mer indices along with their associated read indices and offsets need to be stored in memory— not all the K-mers.

Pairwise comparison of reads roughly runs in $\tilde{\Theta}(N^2L^2)$ if done in the naive way. To speed up the pairwise comparison of noisy reads, one can utilize the fact that the read length is long. We can extract all consecutive f-mers (which act as fingerprints) of the reads and do a lexicographical sort to find candidate neighboring reads and associated offsets for comparison. Since the reads are long, if two reads overlap, there should exist some perfectly matched f-mers which can be identified after the lexicographical sort. This allows an optimized version of Multibridging Algorithm to run in $\tilde{\Theta}(NL \cdot \frac{NL}{G})$ time and $\tilde{\Theta}(NLf)$ space.

X-phased Multibridging

As shown in Fig 1b, when long repeats are flanked approximate repeats, there can be a big gap between the noiseless lower bound and the information requirement for Multibridging Algorithm. A natural question is whether this is due to a fundamental lack of information from the reads or whether Multibridging Algorithm does not utilize all the available information. In this section, we demonstrate that there is an important source of information provided by coverage which is not utilized by Multibridging Algorithm. In particular, we introduce X-phased Multibridging, an assembly algorithm that utilizes the information provided by coverage to phase the polymorphism in long flanked repeat interior. The information requirement of X-phased Multibridging is close to the noiseless lower bound (as shown in Fig 1b) even when some long repeats are flanked approximate repeats.

Description of X-phased Multibridging

Multibridging Algorithm utilizes the random flanking region to differentiate between repeat copies. However, for a flanked approximate repeat, its enclosed exact repeat does not terminate with the random flanking region but only terminates with sparse polymorphism. When we consider the overlap of *two* reads originating from different copies of a flanked approximate repeat, the distinguishing polymorphism is so sparse that it cannot be used to confidently differentiate between repeat copies. Therefore, there is a need to use the extra redundancy introduced by the coverage from *multiple* reads to confidently differentiate between repeat copies and that is what X-phased Multibridging utilizes.

X-phased Multibridging (Alg 3) follows the algorithmic design of Multibridging Algorithm. However, it adds an extra phasing procedure to differentiate between repeat copies of long flanked repeats that Multibridging Algorithm cannot confidently differentiate. We recall that after running step 7 of Multibridging Algorithm, a node in the graph G_{string} corresponds to a substring of the genome and has node label

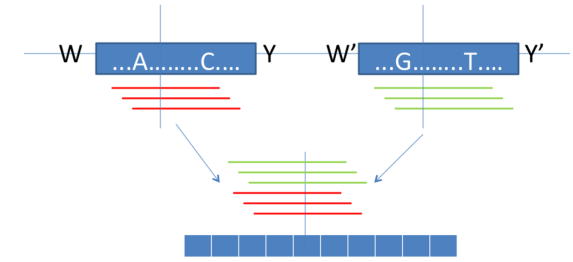

1. Consensus step:
Use all reads to decide where the sites of polymorphism are within the repeat (e.g. Given 60 reads overlap at a base and its counter part in the other copy; If 30 A, 30 G => it is a site of polymorphism; 54 A, 6G => it is not a site of polymorphism)

(a) Consensus Step

2. Read extension step : Use the reads that span the sites of polymorphism to do the vote

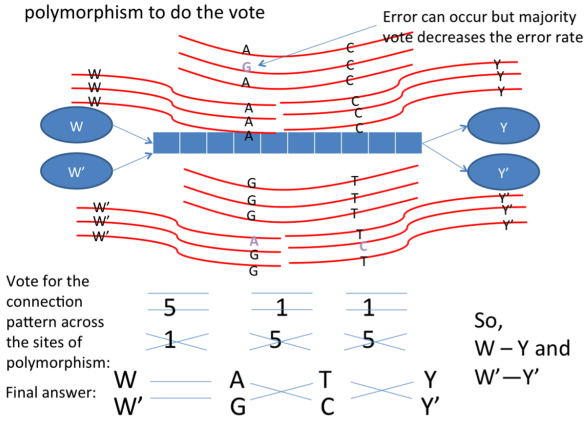

(b) Read Extension Step

Figure 7: Illustration of how to phase polymorphism to extend reads across repeats

consisting of consecutive K-mer cluster indices. An X-node of G_{string} is a node that has in-degree and out-degree ≥ 2. X-node indeed corresponds to a flanked repeat. The incoming/outgoing nodes of the X-node correspond to the incoming/outgoing random flanking region of the flanked repeat.

To be concrete, we focus the discussion on a pair of flanked interleaved repeats, assuming triple repeats are not the bottleneck. However, the ideas presented can be generalized to repeats of more copies.

For the flanked approximate repeat with length $\ell_{\text{int}} < L$ and $\tilde{\ell}_{\text{int}} > L$ (as shown in Fig 7), there is no node-disjoint paths joining incoming/outgoing random flanking region with the distinct repeat copies in G_{string}. It is because the reads are corrupted by noise and the polymorphism is too sparse to differentiate between the repeat copies. Executing Multibridging Algorithm directly will result in the formation of an X-node, which is an artifact due to K-mers from different copies of the flanked approximate repeat erroneously clustered together.

Successful reconstruction requires an algorithm to pair up the correct incoming/outgoing nodes of the X-node(i.e. decide how W, W' and Y, Y' are linked in Fig 7). This is handled by the phasing procedure in X-phased Multibridging, which uses all the reads information. The phasing procedure is composed of two main steps:

- Consensus step: Confidently find out where the sites of polymorphism are located within the flanked repeat interior.
- Read extension step: Confidently determine how to extend reads using the random flanking region and sites of polymorphism as anchors.

Consensus step For the X-node of interest, let D be the set of reads originating from any sites of the associated flanked repeat region and let x_1 and x_2 denote the associated repeat copies. Since the random flanking region is used as anchor, it is treated as the starting base (i.e. $x_1(0) = W$ and $x_2(0) = W'$). For the i^{th} subsequent site of the flanked repeat (where $1 \leq i \leq \tilde{\ell}_{\text{int}}$), we determine the consensus according to Eq (3). This can be implemented by counting the frequency of occurrence of each alphabet overlapping at each site of the repeat. The consensus result determines the sites of polymorphism and the most likely pairs of bases at the sites of polymorphism.

$$\max_{F \subset \{A,C,G,T\}^2} \mathcal{P}(\{x_1(i), x_2(i)\} = F \mid D) \tag{3}$$

Read extension step After knowing the sites of polymorphism, we use those reads that span the sites of polymorphism or random flanking region to help decide how to extend reads across the flanked repeat. Let σ be the possible configuration of alphabets at the sites of polymorphism and random flanking region (e.g. $\sigma = (ACY, GTY')$ means that the two copies of the flanked repeat with the corresponding random flanking region respectively are W–A–C–Y, W'–G–T–Y' where the common bases are omitted).

The following maximum a posteriori estimation is used to decide the correct configuration.

$$\max_{\sigma} \mathcal{P}(\hat{\sigma} = \sigma \mid D, \{x_1(i), x_2(i)\}_{i=1}^{\tilde{\ell}_{\text{int}}}) \tag{4}$$

where $\hat{\sigma}$ is the estimator, D is the raw read set, and x_1, x_2 are the estimates from the consensus step. It is noted that the size of the feasible set for σ is $2^{n_{\text{int}}+1}$.

In practice, for computational efficiency, the maximization in Eq (4) can be approximated accurately even if it is replaced by the simple counting illustrated in Fig 7, which we call count-to-extend algorithm(countAlg). CountAlg uses the raw reads to establish majority vote on how one should extend to the next sites of polymorphism using only the reads that span the sites of polymorphism.

Performance

After introducing the phasing procedure in X-phased Multibridging, we proceed to find its information requirement for successful reconstruction.

The information requirement for X-phased Multibridging is the amount of information required to reduce the error of the phasing procedure to a negligible level.

Algorithm 3 X-phased Multibridging

1. Perform Step 1 to Step 7 of MultiBridging Algorithm
2. For every X-node $x \in G_{string}$
   - a) Align all the relevant reads to the flanked repeat x
   - b) Consensus step: Consensus to find location of polymorphism by solving Eq (3)
   - c) Read extension step: If possible, resolve flanked repeat(i.e. pair up the incoming/outgoing nodes of x) by either countAlg or by solving Eq (4)
3. Perform Step 8 of MultiBridging Algorithm as in Alg 2

p	Coverage (NL/G)	ϵ_1
0.01	20	0.00
0.01	40	0.00
0.01	60	0.00
0.1	20	0.16
0.1	40	0.00
0.1	60	0.00

(a) Calibration for ϵ_1

p	Number of bridging reads k	Upper bound for ϵ_2
0.01	1	0.060
0.01	3	0.0036
0.01	5	0.00024
0.1	11	0.089
0.1	21	0.022
0.1	31	0.0059

(b) Calibration for ϵ_2

Table 1: Calibration of error probability made by the phasing procedure of X-phased Multibridging

The phasing procedure – step 2 in Alg. 3 – is a combination of consensus and read extension steps, which contribute to the error as follows.

Let $\mathcal{E}$ be the error event of the repeat phasing procedure for a repeat, ϵ_1 be the error probability for the consensus step, ϵ_2 be the error probability for the read extension step given k reads spanning each consecutive site of polymorphism within the flanked repeat, δ_{cov} be the probability for having k reads spanning each consecutive sites of polymorphism(i.e. k bridging reads) within the flanked repeat. We have,

$$\mathcal{P}(\mathcal{E}) \leq \epsilon_1 + \epsilon_2 + \delta_{cov} \tag{5}$$

Therefore, to guarantee confidence in the phasing procedure, it suffices to upper bound ϵ_1, ϵ_2 and δ_{cov}. We tabulate the error probabilities of ϵ_1, ϵ_2 in Table 1 for phasing a flanked repeat (whose length is 5000 whereas the genome length is 5M). The flanked repeat has two sites of polymorphism which partition it into three equally spaced segments.

From Table 1, when $p = 0.01$, the information requirement translates to the condition of having three bridging reads spanning the shorter exact repeat of the longest pair of exact interleaved repeats. Therefore, the information requirement for X-phased Multibridging shown in Fig 1b also corresponds to this condition. It is noted that X-phased Multibridging has the same vertical asymptote as the noiseless lower bound. The vertical shift is due to the increase of requirement on the number of bridging reads from $k = 1$ (noiseless case) to $k = 3$ (noisy case).

Simulation of the prototype assembler

Based on the algorithmic design presented, we implement a prototype assembler for automatic genome finishing using reads corrupted by substitution noise. First, the assembler was tested on synthetic genomes, which were generated according to the genome model described previously. This demonstrates a proof-of-concept that one can achieve genome finishing with read length close to ℓ_{crit}, as shown in Fig 8. The number on the line represents the number of simulation rounds (out of 100) in which the reconstructed genome is a single contig with $\geq 99\%$ of its content matching the ground truth.

Second, the assembler was tested using synthetic reads, sampled from genome ground truth downloaded from NCBI. The assembly results are shown in Table 2. The observation from the simulation result is that we can assemble genomes to finishing quality with information requirement near the noiseless lower bound. More information about the detail design of the prototype assembler is presented in the Appendix and source code/data set can be found in [7].

Index	Species	G	p	$\frac{NL}{G}$	L	$\tilde{\ell}_{\max}$	$\tilde{\ell}_{\text{crit}}$	ℓ_{crit}	% match	Ncontig	$\frac{N}{N_{noiseless}}$	$\frac{L}{\ell_{\text{crit}}}$
1	a	1440371	1.5%	37.36 X	930	1817	803	770	100.00	1	1.57	1.21
2	a	1440371	1.5%	33.14 X	970	1817	803	770	99.95	1	1.67	1.26
3	a	1440371	1.5%	29.60 X	1000	1817	803	770	99.99	1	1.66	1.30
4	b	1589953	1.5%	40.82 X	2440	4183	2155	2122	100.00	1	1.30	1.15
5	b	1589953	1.5%	21.31 X	2752	4183	2155	2122	99.99	1	1.19	1.30
6	b	1589953	1.5%	20.66 X	2900	4183	2155	2122	99.99	1	1.35	1.37
7	c	1772693	1.5%	30.03 X	3950	5018	3234	3218	99.96	1	1.36	1.23
8	c	1772693	1.5%	21.96 X	4279	5018	3234	3218	99.97	1	1.33	1.33
9	c	1772693	1.5%	17.03 X	4700	5018	3234	3218	100.00	1	1.31	1.46
10	d	1006701	1.5%	35.23 X	6867	15836	10518	5494	99.05	1	1.72	1.25
11	d	1006701	1.5%	19.88 X	7500	15836	10518	5494	97.86	1	1.30	1.37
12	d	1006701	1.5%	17.69 X	9000	15836	10518	5494	98.10	1	1.68	1.64

Table 2: Simulation results on the assembly of several real genomes using reads corrupted by substitution noise ((a) *Prochlorococcus marinus* (b) *Helicobacter pylori* (c) *Methanococcus maripaludis* (d) *Mycoplasma agalactiae*) with $\ell_{\text{crit}} = \max(\ell_{\text{int}}, \ell_{\text{tri}})$, $\tilde{\ell}_{\text{crit}} = \max(\tilde{\ell}_{\text{int}}, \tilde{\ell}_{\text{tri}})$ and $N_{noiseless}$ is the lower bound on number of reads in the noiseless case for $1 - \epsilon = 95\%$ confidence recovery

Extension to handle indel noise

A further extension of the prototype assembler addresses the case of reads corrupted by indel noise. Similar to the case of substitution noise, tests were performed on synthetic reads sampled from real genomes and synthetic genomes. Simulation results are summarized in Table 3 where p_i, p_d are insertion probability and deletion probability and rate is the number of successful reconstruction(i.e. simulation rounds that show mismatch $< 5\%$) divided by total number of simulation rounds. The simulation result for indel noise corrupted reads shows that X-phased Multibridging can be generalized to assemble indel noise corrupted reads. The information requirement for automated finishing is about a factor of two from the noiseless lower bound for both N and L.

We remark that one non-trivial generalization is the way that we form the noisy De Bruijn graph for K-mer clusters. In particular, we first compute the pairwise overlap alignment among reads, then we use the overlap alignment to group K-mers into clusters. Subsequently, we link successive cluster of K-mers together as

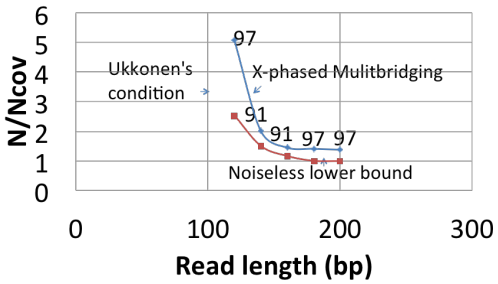

Figure 8: Simulation results on the assembly of synthetic genomes using reads corrupted by substitution noise. The parameters are as follows. $G = 10K; \tilde{\ell}_{\max} = \ell_{\max} = 500, \tilde{\ell}_{\text{int}} = 200, \ell_{\text{int}} = 100$ with two sites of polymorphism within the flanked repeat. $p = 1.5\%, \epsilon = 5\%$.

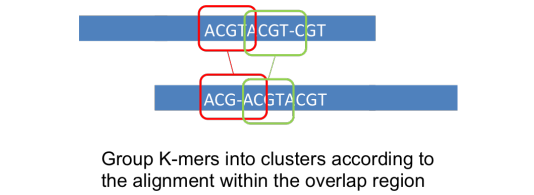

(a) Form K-mer Clusters

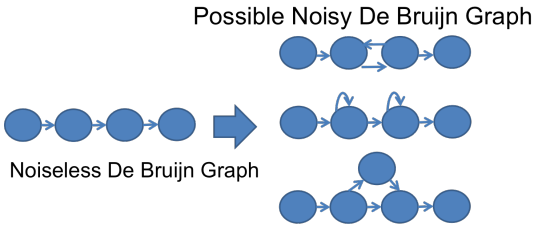

(b) Abnormality in (indel) De Bruijn Graph

Figure 9: Treatment of reads corrupted by indel noise

we do in Alg 2. An illustration is shown in Fig 9a. However, due to the noise being indel in nature, the edges in the noisy De Bruijn graph may point in the wrong direction as shown in Fig 9b. In order to handle this, we traverse the graph and remove such abnormality when they are detected.

Conclusion

In this work, we show that even when there is noise in the reads, one can successfully reconstruct with information requirements close to the noiseless fundamental limit. A new assembly algorithm, X-phased Multibridging, is designed based on a

Type	G	p_i	p_d	$\frac{NL}{G}$	L	$\tilde{\ell}_{\max}$	$\tilde{\ell}_{\text{crit}}$	ℓ_{crit}	$\frac{N}{N_{noiseless}}$	$\frac{L}{\ell_{\text{crit}}}$	Rate
Synthetic	50000	1.5%	1.5%	23.0 X	200	500	200	100	2.25	2	28/30
Synthetic	50000	1.5%	1.5%	24.1 X	180	500	200	100	2.33	1.8	27/30
a	1440371	1.5%	1.5%	28.53 X	1000	1817	803	770	1.60	1.30	1/1
b	1589953	1.5%	1.5%	20.66 X	2900	4183	2155	2122	1.35	1.37	1/1

Table 3: Simulation results on the assembly of real/synthetic genomes using reads corrupted by indel noise(Synthetic: randomly generated to fit $\tilde{\ell}_{\max}, \tilde{\ell}_{\text{crit}}, \ell_{\text{crit}}$; (a) : *Prochlorococcus marinus* ; (b): *Helicobacter pylori*)

probabilistic model of the genome. It is shown through analysis to perform well on the model, and through simulations to perform well on real genomes.

The main conclusion of this work is that, with an appropriately designed assembly algorithm, the information requirement for genome assembly is insensitive to moderate read noise. We believe that the information theoretic insight is useful to guide the design of future assemblers. We hope that these insights allow future assemblers to better leverage the high throughput sequencing read data to provide higher quality assembly.

Competing interests
The authors K.K.L and A.K are or were employees of Pacific Biosciences, a company commercializing DNA sequencing technologies at the time that this work was completed.

Author's contributions
K.K.L, A.K and D.T performed research and wrote the manuscript. K.K.L implemented the algorithms and performed the experiments.

Acknowledgements
The assembly experiments were partly done on the computing infrastructure of Pacific Biosciences.

Declaration
The authors K.K.L and D.T. are partially supported by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

Additional Files
Additional file 1
Details of the proofs, in-depth description of the design of the prototype assembler and details of simulation results are presented.

Author details
[1]Department of Electrical Engineering and Computer Sciences, UC Berkeley, Berkeley, California, United States. [2] Pacific Biosciences, Menlo Park, California, United States.

References
1. Guy Bresler, Ma'ayan Bresler, and David Tse. Optimal assembly for high throughput shotgun sequencing. *BMC Bioinformatics*, 2013.
2. Chen-Shan Chin, David H Alexander, Patrick Marks, Aaron A Klammer, James Drake, Cheryl Heiner, Alicia Clum, Alex Copeland, John Huddleston, Evan E Eichler, et al. Nonhybrid, finished microbial genome assemblies from long-read smrt sequencing data. *Nature methods*, 2013.
3. Sante Gnerre, Iain MacCallum, Dariusz Przybylski, Filipe J Ribeiro, Joshua N Burton, Bruce J Walker, Ted Sharpe, Giles Hall, Terrance P Shea, Sean Sykes, et al. High-quality draft assemblies of mammalian genomes from massively parallel sequence data. *Proceedings of the National Academy of Sciences*, 108(4):1513–1518, 2011.
4. David Gordon, Chris Abajian, and Phil Green. Consed: a graphical tool for sequence finishing. *Genome research*, 8(3):195–202, 1998.
5. Asif Khalak, Ka Kit Lam, Greg Concepcion, and David Tse. Conditions on finishable read sets for automated de novo genome sequencing. *Sequencing, Finishing and Analysis in the Future*, May, 2013.
6. Sergey Koren, Michael C Schatz, Brian P Walenz, Jeffrey Martin, Jason T Howard, Ganeshkumar Ganapathy, Zhong Wang, David A Rasko, W Richard McCombie, Erich D Jarvis, et al. Hybrid error correction and de novo assembly of single-molecule sequencing reads. *Nature biotechnology*, 30(7):693–700, 2012.
7. Ka-Kit Lam, Asif Khalak, and David Tse. www.eecs.berkeley.edu/ ~kakitone.
8. Eric S Lander and Michael S Waterman. Genomic mapping by fingerprinting random clones: a mathematical analysis. *Genomics*, 2(3):231–239, 1988.
9. Elaine Mardis, John McPherson, Robert Martienssen, Richard K Wilson, and W Richard McCombie. What is finished, and why does it matter. *Genome research*, 12(5):669–671, 2002.
10. Duccio Medini, Davide Serruto, Julian Parkhill, David A Relman, Claudio Donati, Richard Moxon, Stanley Falkow, and Rino Rappuoli. Microbiology in the post-genomic era. *Nature Reviews Microbiology*, 6(6):419–430, 2008.
11. Abolfazl Motahari, Kannan Ramchandran, David Tse, and Nan Ma. Optimal dna shotgun sequencing: Noisy reads are as good as noiseless reads. *Proceedings of the 2013 IEEE International Symposium on Information Theory, Istanbul, Turkey, July 7-12, 2013*, 2013.
12. Giuseppe Narzisi and Bud Mishra. Comparing de novo genome assembly: The long and short of it. *PLoS ONE*, 6(4):e19175, 04 2011.
13. Mihai Pop. Genome assembly reborn: recent computational challenges. *Briefings in bioinformatics*, 10(4):354–366, 2009.
14. DNA SEQUENCING. A plan to capture human diversity in 1000 genomes. *Science*, 21:1842, 2007.
15. Claude E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423, 623–656, July, October 1948.
16. Jared T Simpson and Richard Durbin. Efficient de novo assembly of large genomes using compressed data structures. *Genome Research*, 22(3):549–556, 2012.
17. Peter J Turnbaugh, Ruth E Ley, Micah Hamady, Claire M Fraser-Liggett, Rob Knight, and Jeffrey I Gordon. The human microbiome project. *Nature*, 449(7164):804–810, 2007.
18. Esko Ukkonen. Approximate string-matching with q-grams and maximal matches. *Theoretical computer science*, 92(1):191–211, 1992.
19. Daniel R Zerbino and Ewan Birney. Velvet: algorithms for de novo short read assembly using de bruijn graphs. *Genome research*, 18(5):821–829, 2008.

Appendix : Proof on Performance Guarantee

Here we use the short hand of l_{repeat}, $l_{interleaved}$ and l_{triple} to represent the corresponding approximate repeat length of longest simple, interleaved, triple repeat respectively.

Greedy Algorithm

Let us define a $\theta-$ neighborhood of the repeat specified by $x[a:b]$ and $x[c:d]$ to be the loci of repeat which are $x[a-\theta:b+\theta]$ and $x[c-\theta:d+\theta]$.

We say a repeat is $\theta-$bridged if there exists a read that cover the θ-neighborhood of at least one copy of the repeat. For simplicity of arguments, we assume $l_{repeat} >> \max(l_{intereleave}, l_{triple})$.

Lemma 3 *We first note the following sufficient conditions for Noisy Greedy to succeed.*

1. *Merging at stages from L to $\ell_{\text{iid}}(p,\epsilon,G)$ are merging successive reads*
2. *Every successive reads have overlap with length at least $\ell_{\text{iid}}(p,\frac{\epsilon}{3},G)$*

Theorem 4 *Under the generative model on genome, with $\ell_{\text{iid}} = \ell_{\text{iid}}(p,\frac{\epsilon}{3},G)$, $\alpha = \alpha(p,\frac{\epsilon}{3},G)$, if*

$$L > l_{repeat} + 2\cdot\ell_{\text{iid}}$$
$$G > N > \max\left(\frac{G\ln\frac{3}{\epsilon}}{L - l_{repeat} - 2\cdot\ell_{\text{iid}}}, \frac{G\cdot\ln\frac{N}{\epsilon/3}}{L-\ell_{\text{iid}}}\right)$$

, then $\mathcal{P}(\mathcal{S}^C) \le \epsilon$.

Proof In order to prove that claim, let us break down into several subparts

Let E_1 be the event that condition 1 in Lemma (3) is not satisfied. E_2 be the event that condition 2 in Lemma (3) is not satisfied. E_3 be the event that the long/interleave/triple repeat is not $\ell_{\text{iid}}-$bridged.

Now we claim that with the chose (N, L) in the range,

1. $\mathcal{P}(E_1) \le \frac{\epsilon}{3}$
2. $\mathcal{P}(E_3) \le \frac{\epsilon}{3}$
3. $\mathcal{P}(E_2 \mid E_1^C \cap E_3^C) \le \frac{\epsilon}{3}$

We first see how we can use these to obtain the desired claim and proceed to prove each of the above sub-claims.

$$\begin{aligned}
\mathcal{P}(S^C) =& \mathcal{P}(E_1) + \mathcal{P}(E_2 \cap E_1^C) \\
=& \mathcal{P}(E_1) + \mathcal{P}(E_2 \cap E_1^C \cap E_3^C) + \mathcal{P}(E_2 \cap E_1^C \cap E_3) \\
\le& \mathcal{P}(E_1) + \mathcal{P}(E_2 \mid E_1^C \cap E_3^C) + \mathcal{P}(E_3) \\
\le& \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \\
=& \epsilon
\end{aligned}$$

Now, we proceed to prove each of the sub-claims.

1. With $N > \frac{G\cdot\ln\frac{N}{\epsilon/3}}{L-\ell_{\text{iid}}}$, we have,

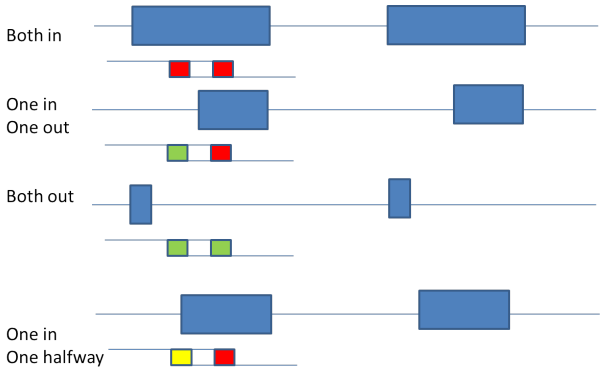

Figure 10: Overlap Type

$$\begin{aligned}\mathcal{P}(E_1) \leq & N\exp(-\frac{N}{G}(L-\ell_{\text{iid}})) \\ \leq & \frac{\epsilon}{3}\end{aligned}$$

2. With $N > \frac{G\cdot\ln\frac{3}{\epsilon}}{L-l_{repeat}-2\cdot\ell_{\text{iid}}}$ we have,

$$\begin{aligned}\mathcal{P}(E_3) \leq & \exp(-\frac{N}{G}(L-2\ell_{\text{iid}}-l_{repeat})) \\ \leq & \frac{\epsilon}{3}\end{aligned}$$

3. With the choice of $\ell_{\text{iid}} = \ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$, we have,

$$\begin{aligned}\mathcal{P}(E_2 \mid E_1^C \cap E_3^C) \leq & N^2\cdot\exp(-\ell_{\text{iid}}D(\alpha||\frac{3}{4})) \\ & + 2N\exp(-\ell_{\text{iid}}D(\alpha||\eta)) \\ \leq & G^2\cdot\exp(-\ell_{\text{iid}}D(\alpha||\frac{3}{4})) \\ & + 2G\exp(-\ell_{\text{iid}}D(\alpha||\eta)) \\ \leq & \frac{\epsilon}{3}\end{aligned}$$

Here we use the fact that there are indeed 4 types of overlap as in Fig 10. And given the bridging condition, we are only left with 2 types, namely, both ending segments outside/exactly one ending segment outside the longest repeat repeat region.

Simple De Bruijn Algorithm

Before continuing proving the performance of Multibridging Algorithm, it is instructive to analyze the following Simple De Bruijn Algorithm (Alg 4) because this is closely related to the Multibridging Algorithm.

Algorithm 4 Noisy Simple De Bruijn

0. Choose K to be $\max(l_{interleave}, l_{triple})$
1. Extract Kmers from reads
2. Clusters Kmers
3. Form Kmer Graphs
4. Condense the graph
5. Clear Branches
6. Condense graph
7. Find Euler Cycle

Here we first define several genomic region of interest which we will refer to in the proofs below.

a) S_0 = set of K-mers that are completely inside $\ell_{\text{iid}}-$ neighborhood of the longest repeat

b) S_1 = set of K-mers that are completely inside the longest repeat

c) $S_2 = S_0 \backslash S_1$

Lemma 5 *Here we provide several deterministic conditions that guarantee the success of the algorithm.*

1. Successive reads overlap with length at least K

2 K-mers are almost correctly clustered, that is,

a) K-mers from the same locus but not merged

b) x not in S_0 s.t. x get clustered with wrong K-mers

c) x in S_0 s.t.x get clustered with elements other than its own cluster/mirror cluster(mirror cluster is defined to be the cluster for the other copy of the repeat)

3) Repeat at both circle are at least $2 \cdot \ell_{\text{iid}}$ *separated(the interleaving segments between the repeat differ with at least* $2\ell_{\text{iid}}$ *in length)*

Proof We note that every length K segments $x \notin S_0$, they are represented as a distinct node in the K-mer graph because of the length K that we pick and the condition that successive reads overlap at least K bases. Moreover, for K-mers $x \in S_1$, they are condensed into the repeat as 'X' in Fig 11a. However, for the K-mers $x \in S_2$, they have chances not to merge properly, thus they form into the branches surrounding 'X' in Fig 11a. Because of condition 3, branch clearing will not eliminate the 'A' or 'C' in Fig 11a, further after condensing, we get the desired K-mer graph as in Fig 11b and this can be successfully read by a Eulerian Walk.

Theorem 6 *If* $G > \frac{6}{\epsilon\ell_{\text{iid}}}, G \geq N \geq \frac{G \cdot \ln \frac{3N}{\epsilon}}{L - \max(l_{int}, l_{triple}) - 2\ell_{\text{iid}}}$ *, with* $\ell_{\text{iid}} = \ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$ $\alpha = \alpha(p, \frac{\epsilon}{3}, G)$*, then,* $\mathcal{P}(S^C) \leq \epsilon$

Proof We first note that in order to obtain a bound on the error probability, we only need to separately bound the probability that each of the conditions in Lemma 5 fail,which are $\leq \frac{\epsilon}{3}$ each. Thus, combining, we get, $\mathcal{P}(S^C) \leq \epsilon$.

Multibridging Algorithm

An illustration of noisy Multibridging Algorithm is shown in Fig (12).

Lemma 7 *Here are the deterministic conditions for the algorithm to succeed.*

1) Every successive reads overlap at least $\ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$

2) K-mers are almost correctly clustered, that is,

a) K-mers from the same locus but not merged

b) x not in S_0 s.t. x get clustered with wrong K-mers

c) x in S_0 s.t. x get clustered with elements other than its own cluster/mirror cluster

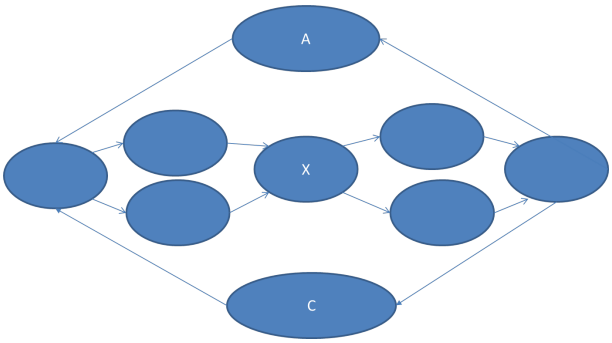

(a) Before branch clearing

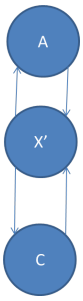

(b) After branch clearing

Figure 11: Branch clearing

3) Repeat at both circle are at least $2 \cdot \ell_{\text{iid}}$ *separated*

4) When finding successors/predecessors, they are the real successors and predecessors

Proof Along the same lines as the proof in Lemma (5), we only note that in this algorithm, we have an extra step of finding predecessor/successors. Moreover, the overlap here is significantly reduced to only ℓ_{iid} instead of K in the Noisy Simple De Bruijn case.

Theorem 8 *With* $G > \frac{6}{\epsilon \ell_{\text{iid}}}$, $G \geq N \geq \max(\frac{G}{L-2\ell_{\text{iid}}} \ln \frac{N}{\epsilon/3}, \frac{G \ln \frac{3}{\epsilon}}{L-\max(l_{triple}, l_{interleave})-2\ell_{\text{iid}}})$ *,with* $\ell_{\text{iid}} = \ell_{\text{iid}}(p, \frac{\epsilon}{3}, G)$ $\alpha = \alpha(p, \frac{\epsilon}{3}, G)$ *, then* $\mathcal{P}(S^C) \leq \epsilon$.

Proof Here we note that with the given coverage, bridging conditions of the interleave repeat and the triple repeat are satisfied. And when this is true, then Condition 4 in Lemma 7 is true with high probability. Following the arguments in Theorem 6, we get desired.

Appendix: Design and additional algorithmic components for the prototype assembler

Pipeline of the prototype assembler

The pipeline of the prototype assembler is shown in Fig 13. With a ground truth genome as input, the output is the performance of the whole pipeline by giving the mismatch rate.

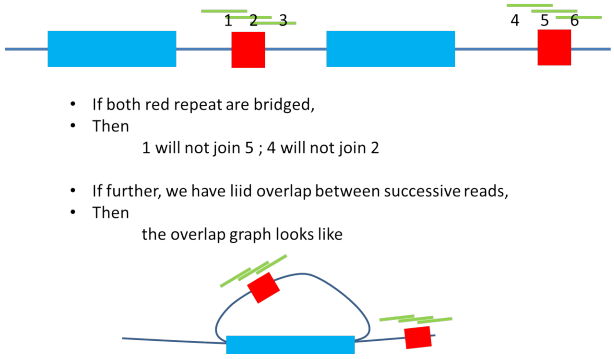
1 2 3
4 5 6
If both red repeat are bridged,
Then
1 will not join 5 ; 4 will not join 2
If further, we have liid overlap between successive reads,
Then
the overlap graph looks like

Figure 12: An illustration of the Noisy Multi Bridging

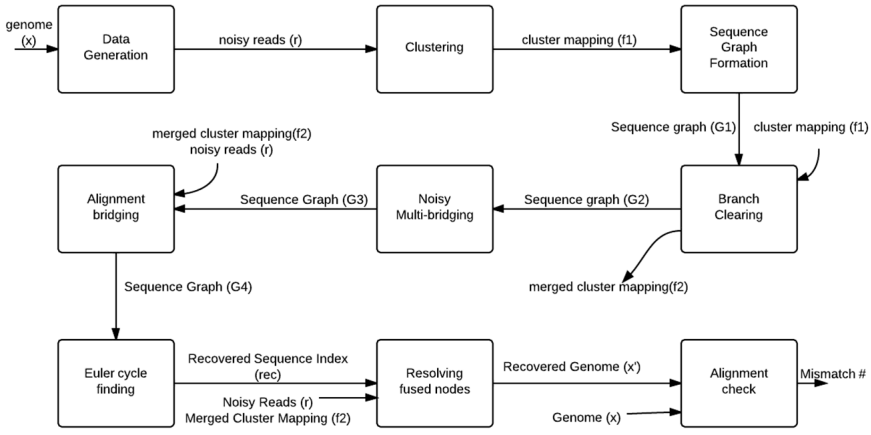

Figure 13: Pipeline of the prototype assembler

A more robust branch clearing step

Since we employ a speed up step in the clustering and there may be K-mers that are not completely clustered correctly in the clustering step of Multibridging Algorithm. Regarding that, we need to have a more robust branch clearing step. In particular, we first classfy nodes as "big" or "small" nodes based on the size of the nodes in the sequence graph. The key idea is to merge the small nodes together while keeping the big nodes unchanged. Starting from each big nodes, we tranverse the graph to detect all the small nodes that link the current big node to other big nodes. Then, we classify the small nodes into levels(depending on its distance from the current big node). After that, the small nodes in the same level are merged. Finally, we note that we keep the reachability among each big nodes.

Enhanced Multibridging Algorithm that can resolve middle range repeats

We note that the ideas presented here can also be found in the prior work on the treatment of noiseles sreads. It is stated here for completeness. In the noisy setting, instead of considering the alphabet set to be $\Sigma = \{A, C, G, T\}$, one can consider the alphabet set as the cluster index of the K-mers.

Algorithm 5 Enhanced Multibridging Algorithm

Resolution of repeats:
0. Intially the weight of the edge are set to be 1.
1. While there is a X-node v :
a) For each edge (p_i,v) with weight $a_{p_i,v}$, create a new node $u_i = {}^{p_i}\!\rightarrow v$ and an edge (p_i, u_i) with weight $1 + a_{p_i v}$. Similarly, for each edge (v,q_j), create a new node $w_j = v^{\rightarrow q_j}$ and an edge (w_j, q_j)
b) If v has a self-loop (v, v) with weight $a_{v,v}$, add an edge $({}^{v\rightarrow}v, v^{\rightarrow})$ with weight $a_{v,v} + 2$
c) Remove node v and all incident edges
d) For each pair of u_i, w_j adjacent in a read(extending to at least length of ℓ_{iid} on both sides of the X-node), add edge (u_i, w_j). If exactly on each of the $u_i$and $w_j$nodes have no added edge, add the edge.
e) Condense the graph

Appendix: Treatment of indel noise

Formation of K-mer De Bruijn graph for indel corrupted reads

In order to form K-mer De Bruijn graph for indel corrupted reads, we first need to have a clear notion of K-mers. We define K-mers to be the length K segments in the genome ground truth (as opposed to the usual definition from the reads). Although we mostly work on the reads themselves, the definition of the Kmers are based on the ground truth. In order to successfully cluster K-mers, we need to do the following steps.

1. We first compute the pairwise alignement of the reads.

2. Based on the pairwise alignment, for each length K-segments, we know which should be aligned to which. We then group them together using the alignment result.

3. Finally, we end up with the length K segments from the reads clustering together, and now we use it as an operational way to identify the Kmers since each cluster will naturally correspond to a K-mers originated from the genome groundtruth(though there are a few discrepancy, mostly this is correct).

4. After we identify the K-mers clusters, we add an edge between them if there exists a read such that there are two consecutive Kmers originate from it.

Graph surgery to clear abnormality of the noisy De Bruijn graph

Due to indel noise and runs of the same alphabet, the way that we form K-mers graph may need to abnormality of the graph. We thus perform a graph tranversal and identify the abnormality that are of short length(i.e. resulted from noise but not the genome structure). After that, we remove such abnormality. This step also involves transitive edge reduction and removal of small self loops.

X-phased step tailored for indel noise type

Generalization to handle Indel Error

When dealing with indel noise, the neighborhood of reads can also affect consensus of the base. There we have to do sequence alignment in order to find the appropriate posterior probability in order to do a maximum likelihood estimate of whether a particular given genomic location is a site of polymorphism or not. In order to do that, we formulate the problem as a ML problem as follows.

Repeat Type	C_s	L_s	C_l	L_l	p_{del}	p_{ins}	G	Homology	l^{approx}	l^{exact}	Success %
Randomly generated	50X	100	50X	240	10%	10%	10000	0.67%	300	150	99%
A repeat of Ecoli-K12	–	–	80X	3000	10%	10%	4646332	0.48%	5182	1507	89%
A repeat of Bacillus anthracis	–	–	80X	3500	10%	10%	5227293	0.23%	4778	2305	85%
A repeat of Meiothermus ruber	–	–	80X	750	10%	10%	3097457	1.40%	1217	257	94%

Table 4: Simulation results on long contig creator(C_s, L_s are coverage and readlength for short reads. C_l, L_l are coverage and readlength for long reads. p_{del}, p_{ins} are the probability of insertion and deletion. G is the length of the genome. Homology is the number of SNPs divided by the length of the approximate repeat. l^{approx}, l^{exact} are the length of the approximate and exact repeat being studied. Success % is the percentage of success in 100 rounds)

$$\max_{T\in\Omega} \quad \Pi_{i\in S} P(R_i \mid T), \qquad P_{err} = P_{opt} \tag{6}$$

$$\max_{T\in\Omega'} \quad \Pi_{i\in S} P(R_i \mid T), \qquad P_{err} = P_{opt} + \delta_1 \tag{7}$$

$$\max_{T\in\Omega'} \quad \Pi_{i\in S'} P(R_i \mid T), \qquad P_{err} = P_{opt} + \delta_1 \tag{8}$$

$$\max_{T\in\Omega'} \quad \Pi_{(j,k)}\Pi_{i\in S'_{jk}} P(R_i \mid T_j^{j+k}), \qquad P_{err} = P_{opt} + \delta_1 \tag{9}$$

$$\max_{T\in\Omega'} \quad \Pi_{(j,j+1)}\Pi_{i\in S'_{j,j+1}} P(R_i \mid T_j^{j+1}), \qquad P_{err} = P_{opt} + \delta_1 + \delta_2 \tag{10}$$

Here we also discuss about the places that we take approximation to enhance the computational efficiency in the steps of the previous reduction. From (6) to (7), we use some heuristics to find out the possible location of SNPs within the whole repeat in which disagreement is observed after several rounds of error correction. From (7) to (8), we remove all the reads that only span one single SNPs and it has no effect on the error of the detection problem that we are trying to solve. From (8) to (9), we further partition the reads into group in which S'_{jk} is the set of reads that only span the SNPs j to j+k. Doing this can decompose the ML problem into smaller subproblems with no effect on the accuracy. Finally, in practice, we take a first order approximation of (9) to (10) by only onsidering two SNPs for each subproblem.

As for each of the marginal probability distribution, the best way is to run Sum-Product algorithm to compute in a dynamic programming fashion similar to S-W alignment. But as pointed out in Quiver, this steps can be significanly speeded up using a Viterbi approximation and this is also what we implemented in the simulation code.

Simulation study

We simulated on both synthetic and real data set with indel noise and on a double stranded DNA. In the simulation, we assume that the reads from the neighborhood of a repeat is given and our goal is to decide how to extend the reads to span the repeat copies into the flanking region correctly. The correctness is evaluated based on whether they can correctly extend the correct reads into the flanking region.

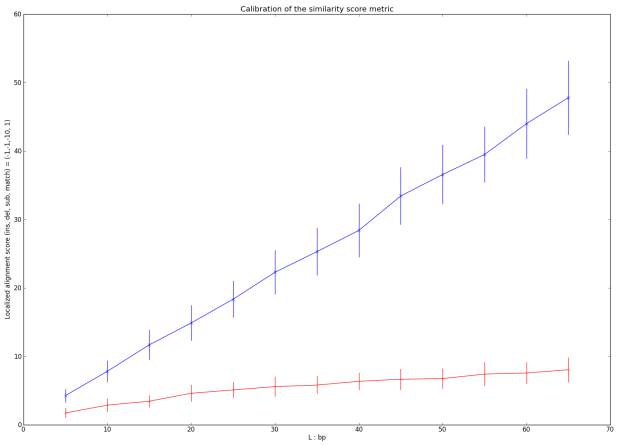

Figure 14: A calibration for similarity score using global alignment computation.

Edit distance metric calibration

We also do a study on whether we can use alignment score to differentiate between segments from being extracted from the same locus or not. In Fig 14, the upper curve is the score for segment extracted from the same locus while the bottom curve is for completely iid randomly(irrelevant) generated segment. And we simulate it for 100 times at each length and the bar indicate 1 standard deviation from the mean.

Tolerance in the Multibridging step

We note that due to the indel noise and the graph surgery that we perform, an X-node of the graph may be p times longer than the usual size of the approximate repeat, thus we should have a corresponding higher tolerance to use the reads to bridge across the repeats.

Computation speed up of alignment step

The key bottleneck in computation speed of the indel extension is on the pairwise alignment of the reads, which can be speeded up using the ideas in BLAST. We use sorting to identify exact matching fingerprint that identify the starting and ending location of the segment that need to be aligned with. After that, we do a local search instead of the whole dynamic programming search.

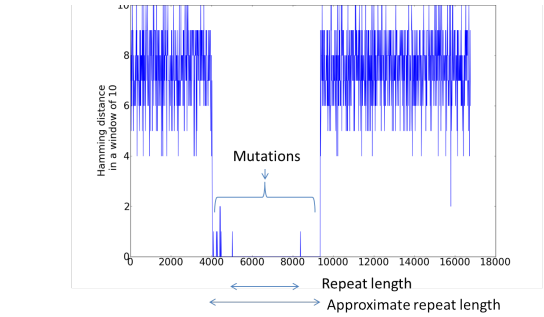
Hamming distance in a window of 10
Mutations
Repeat length
Approximate repeat length
Stretch = Approximate repeat length/ Repeat length;
mutation rate = # of mutation / (Approximate repeat length – Repeat length)

Figure 15: Example of how to define stretch and mutation rate

Appendix : Evidence behind model

Approximate Repeat

We let the underlying genome be $\vec{x}$ and use the short hand that $x[a:b]$ be the a^{th} to $(b-1)^{th}$ entries of $\vec{x}$.

Let $\vec{v}_1 = x[s_1 : s_1 + l]$ and $\vec{v}_2 = x[s_2 : s_2 + l]$ be two length l substrings of the genome with starting positions at s_1 and s_2 respectively. We call $\vec{v}_1$ and $\vec{v}_2$ be an approximate repeat of length l if

$d(x[s_1 - W : s_1], x[s_2 - W : s_2]) \geq 0.7W$

$d(x[s_1 + l : s_1 + l + W], x[s_2 : s_2 + l + W]) \geq 0.7W$

$d(x[s_1 - W + k : s_1 + k], x[s_2 - W + k : s_2 + k]) < 0.7W$ for all $0 < k < l$

To understand approximate repeat better, we plot the Hamming distance for consecutive disjoint window of length 10 as shown in Fig 15 .

Classification of approximate repeat

While repeats are studied in the literature[**?**], they are not investigated by looking at the ground truth. This is partially due to the insufficiency of data in the early days of genome assembly development. Therefore, based on the ground truth genome, we define several quantities that allow us to classify approximate repeat and understand the approximate repeat spectrum of genome. Here we define stretch and mutation rate. Stretch is defined to be the ratio of the length (l^*) of the longest exact repeat within an approximate repeat divided by the length (l_{approx}) of the approximate repeat. Mutation rate is defined to be number of mutation within approximate repeat divided by ($l_{approx} - l^*$). An illustration is shown in Fig 15.

Moreover, we do a scatter plot to classify the approximate repeats(approximate repeat having exact repeat length within top 20) and we have a plot of approximate repeat spectrum as in Fig 16.

From the plots in Fig 16, we classify approximate repeat as homologous repeat if the stretch is bigger than 1.25 and as non-homologous repeat if the stretch is less than 1.25.

For the scatter plot, every approximate repeat is a dot there with x coordinate and y coordinate being mutation rate and stretch respectively. And the color represents the length of that approximate repeat. For the approximate repeat spectrum

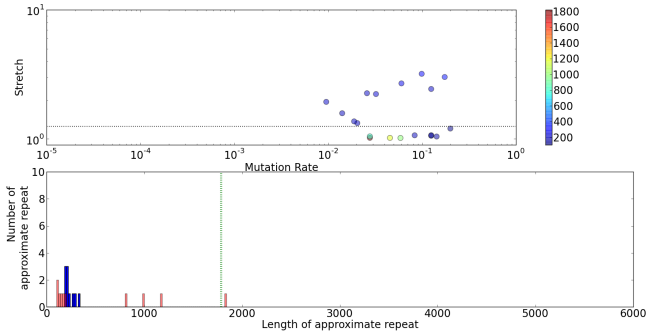

Figure 16: Classification of approximate repeats and approximate repeat spectrum. The upper plot is scatter plot to classify approximate repeat. The lower plot is the approximate repeat spectrum

plot, the red bar represent non-homogeneous repeat while the blue bar represent homologous repeat. The green dotted line indicates the length of the longest repeat.

We focus on genomes when the non-homologous repeat dominates, namely the longest interleave and the longest triple repeats are non-homologous because the stretch is relatively short which can be captured by our generative model. We do not distinguish between the length of approximate or exact repeat are considered to be the same and we do not distinguish between the two in the discussion because of the small stretch.

Stopping criterion for defining approximate repeat by MLE estimate

Parametric model

Let L_k be the number of bases between the $(k-1)^{th}$ and the k^{th} SNPs starting from the right end-point of a repeat.

We consider the following probabilistic model for the L_k. $\{L_k\}_{k=1}^n$is taken as an independent sequence of geometrically distributed random variables with parameter $\Theta = \{p_1, p_2, r\}$ defined as follows.

$$L_k \sim \begin{cases} Geo(p_1) & \text{if } 1 \leq k \leq r \\ Geo(p_2) & \text{if } r < k \leq n \end{cases}$$

MLE estimate of parameters

We now would like to estimate Θ given the observation of $\{\hat{L}_k\}_{k=1}^n$ by maximum likelihood estimation. Consider the log-likelihood function $L(\Theta) = \log \mathcal{P}(\{\hat{L_k}\}_{k=1}^n \mid \Theta)$.

$$\begin{aligned} L(\Theta) &= \log \mathcal{P}(\{\hat{L_k}\}_{k=1}^n \mid \Theta) && (11) \\ &= \log\{[\Pi_{k=1}^r (1-p_1)^{\hat{L_k}} p_1] \cdot [\Pi_{k=r+1}^n (1-p_2)^{\hat{L_k}} p_2]\} && (12) \\ &= r \log p_1 + [\sum_{k=1}^r \hat{L}_k] \cdot \log(1-p_1) + (n-r) \cdot \log p_2 + [\sum_{k=r+1}^n \hat{L}_k] \cdot \log(1-p_2) && (13) \end{aligned}$$

Algorithm 6 Linear time algorithm to estimate the stopping criterion

1.
a)$C_0 \leftarrow 0$
b) for r = 1 to n
$C_r \leftarrow C_r + \hat{L}_r$
2.
a)$D_0 \leftarrow C_n$
b)for r = 1 to n
$D_r \leftarrow C_n - \hat{L}_r$
3. for r = 1 to n
$\hat{p}_1^{(r)} \leftarrow \frac{1}{1+\frac{C_r}{r}}$
$\hat{p}_2^{(r)} \leftarrow \frac{1}{1+\frac{D_r}{n-r}}$
$\Theta_r \leftarrow (r, \hat{p}_1^{(r)}, \hat{p}_2^{(r)})$
$X_r \leftarrow L(\Theta_r)$
4. find maximum among $\{X_r\}_{r=1}^n$, and the corresponding Θ_r is the MLE estimate.
5. (Differentiate between homologous and non-homologous repeat)
If the optimal $\hat{p}_1^{(r)}, \hat{p}_2^{(r)}$are too close (i.e. $\hat{p}_1^{(r)} > 0.2$), then claim $r = 1$; else, claim $\hat{r} = r$.

And we want to find $\hat{\Theta} = \arg\max_{\Theta} L(\Theta)$.

Observe that, if we fix $1 \leq r \leq n$, then the optimal $\hat{p}_1$ and $\hat{p}_2$ can be readily obtained by taking derivative on $L(\Theta)$ with respect to $p_1$and p_2, specifically,

$$\hat{p_1} = \frac{1}{1 + \frac{\sum_{k=1}^{r} \hat{L}_k}{r}} \tag{14}$$

$$\hat{p}_2 = \frac{1}{1 + \frac{\sum_{k=r+1}^{n} \hat{L}_k}{n-r}} \tag{15}$$

$\hat{\Theta}$ can then be obtained by running over all integral $1 \leq r \leq n$ and use the corresponding optimal $\hat{p}_1$ and $\hat{p}_2$ to obtain the $L(\Theta)$, and finally we use the r that gives the highest value of $L(\Theta)$ as the MLE estimate given the observation.

Linear time algorithm to estimate the stopping criterion

Moreover, this can be done by the following algorithm Algo 6, which run in linear time $\Theta(n)$ with respect to the number of observations n.

A sample plot is of who we can use the critierion to accurately define the ending of approximate repeat is shown in Fig 17.

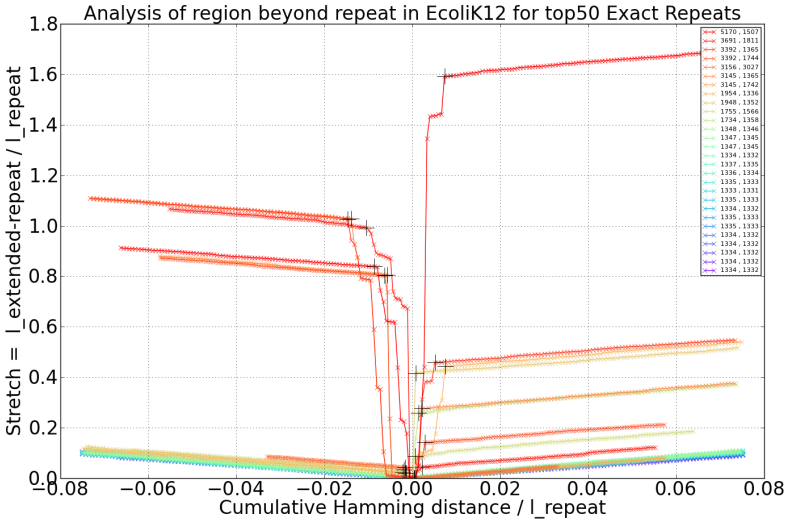

Figure 17: An example plot that define the stopping point of approximate repeat by Algorithm 6

Appendix of the dot plot of finished genomes

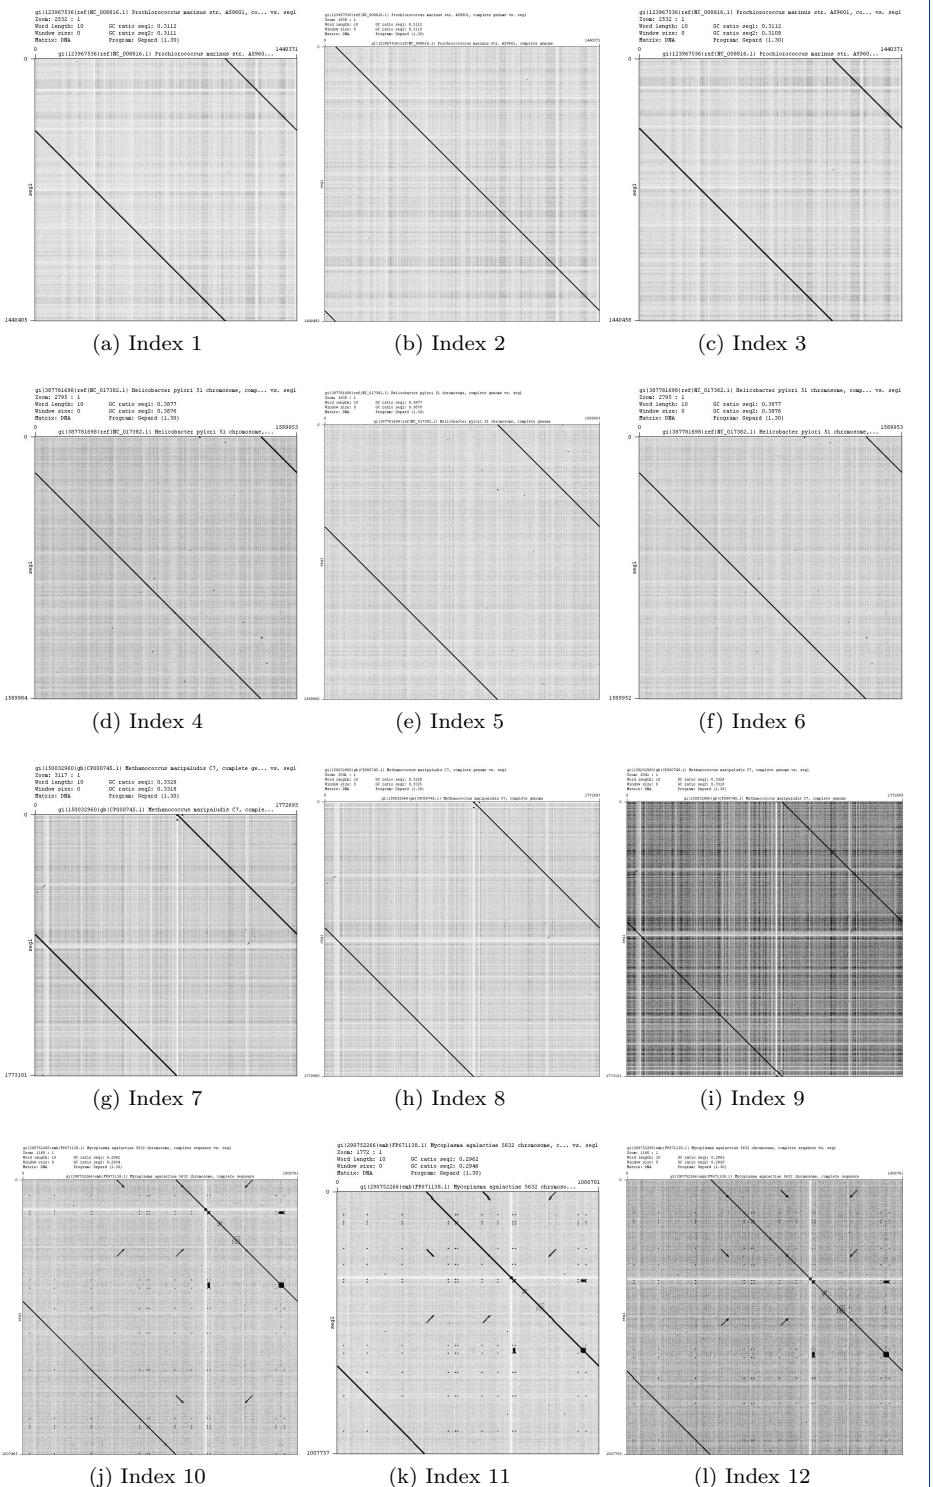

(a) Index 1 (b) Index 2 (c) Index 3

(d) Index 4 (e) Index 5 (f) Index 6

(g) Index 7 (h) Index 8 (i) Index 9

(j) Index 10 (k) Index 11 (l) Index 12

Figure 18: Dot plot of recovered genomes against ground truth(according to index in Table 2)